CMSC 611: Advanced Computer Architecture

Memory & Storage

Some material adapted from Mohamed Younis, UMBC CMSC 611 Spr 2003 course slides

Main Memory Background

- Performance of Main Memory:
  - Latency: affects cache miss penalty
    - Access Time: time between request and word arrives
    - Cycle Time: time between requests
  - Bandwidth: primary concern for I/O & large block
- Main Memory is DRAM: Dynamic RAM
  - Dynamic since needs to be refreshed periodically
  - Addresses divided into 2 halves (Row/Column)
- Cache uses SRAM: Static RAM
  - No refresh
    - 6 transistors/bit vs. 1 transistor/bit, 10X area
  - Address not divided: Full address

DRAM Logical Organization

4 Mbit DRAM:

square root of bits per RAS/CAS

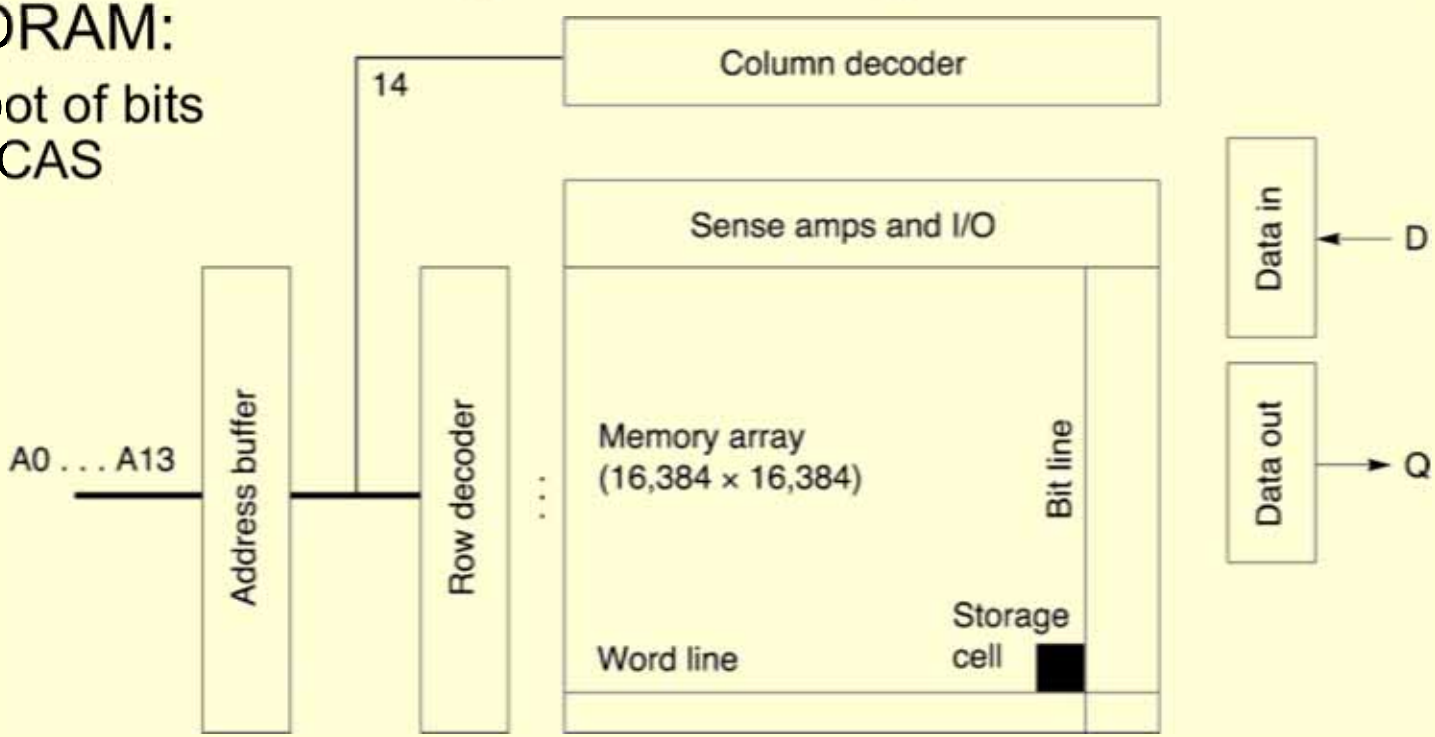

- Refreshing prevent access to the DRAM (typically 1-5% of the time)
- Reading one byte refreshes the entire row
- Read is destructive and thus data need to be re-written after reading
  - Cycle time is significantly larger than access time

Processor-Memory Performance

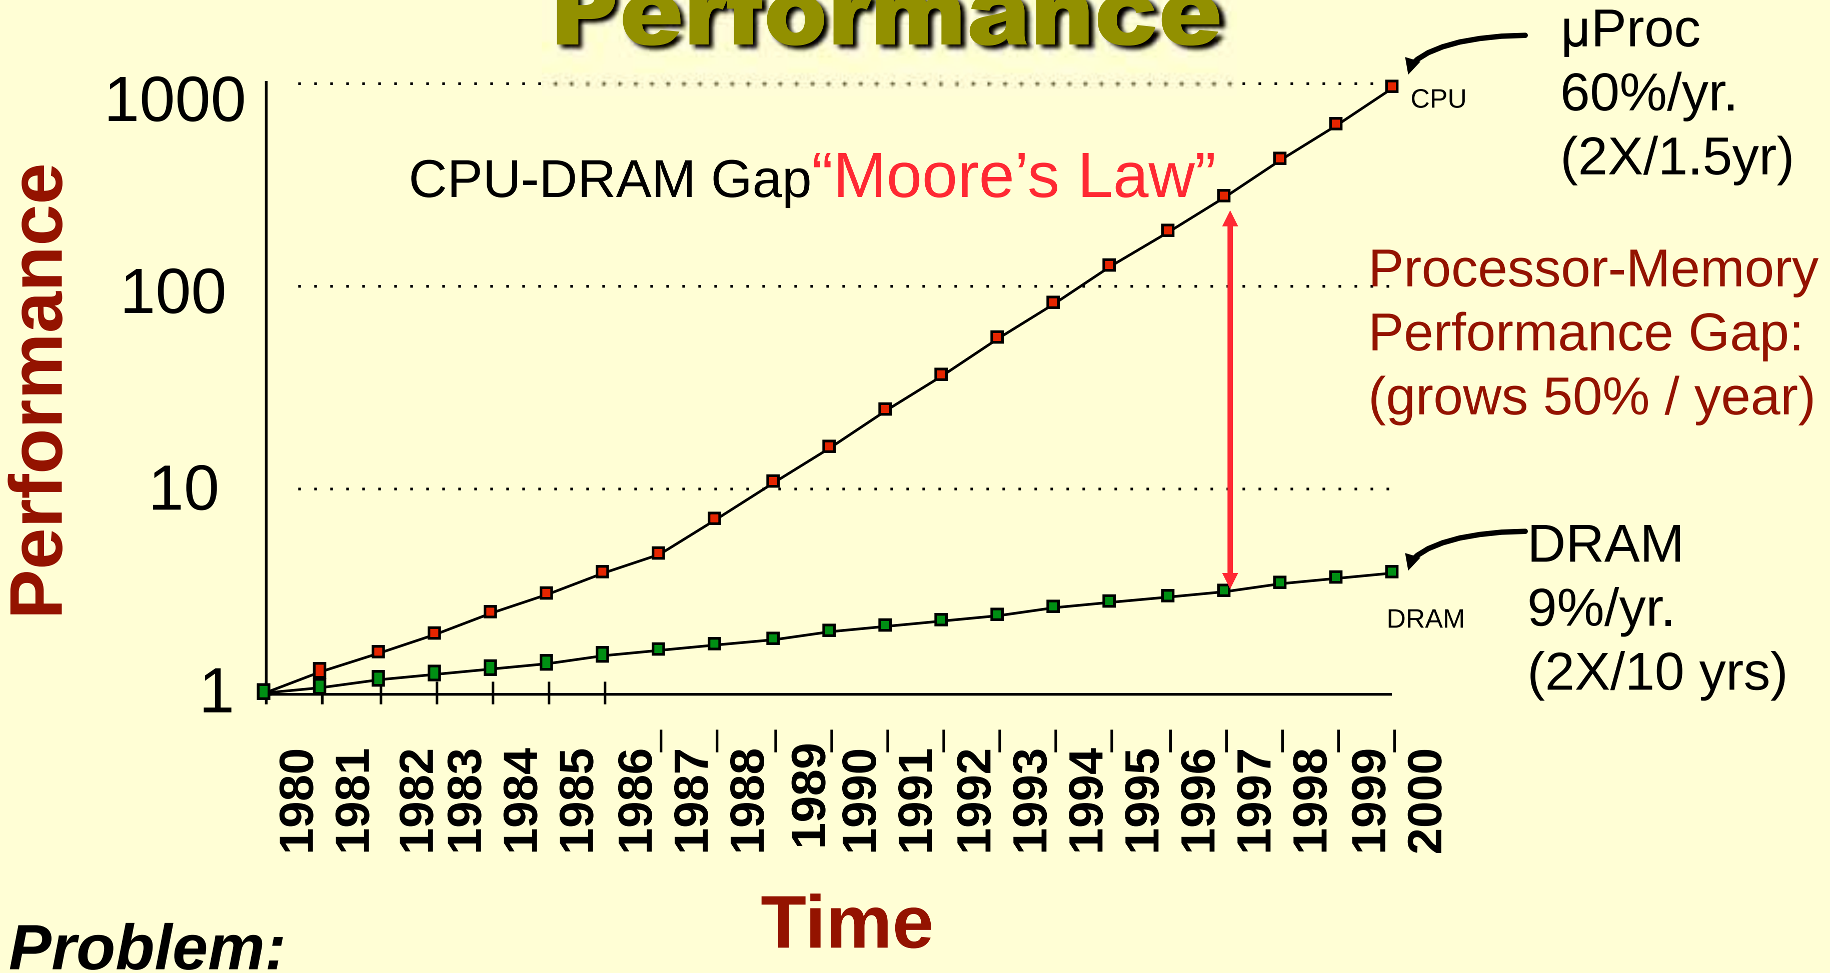

Problem:

Improvements in access time are not enough to catch up

Solution:

Increase the bandwidth of main memory (improve throughput)

Memory Organization

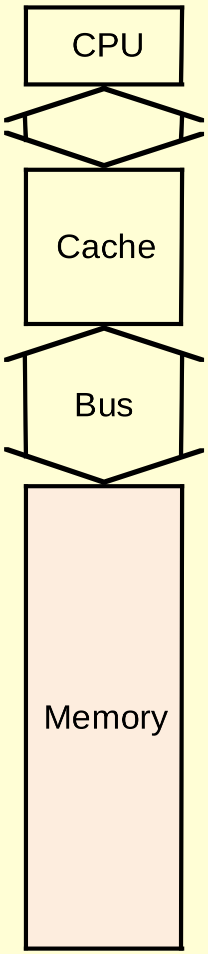

a. One-word-wide memory organization

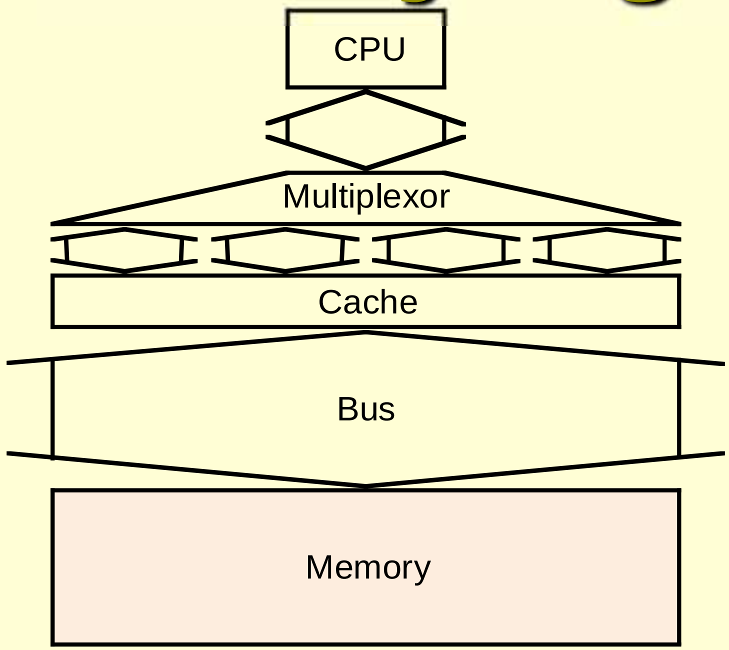

b. Wide memory organization

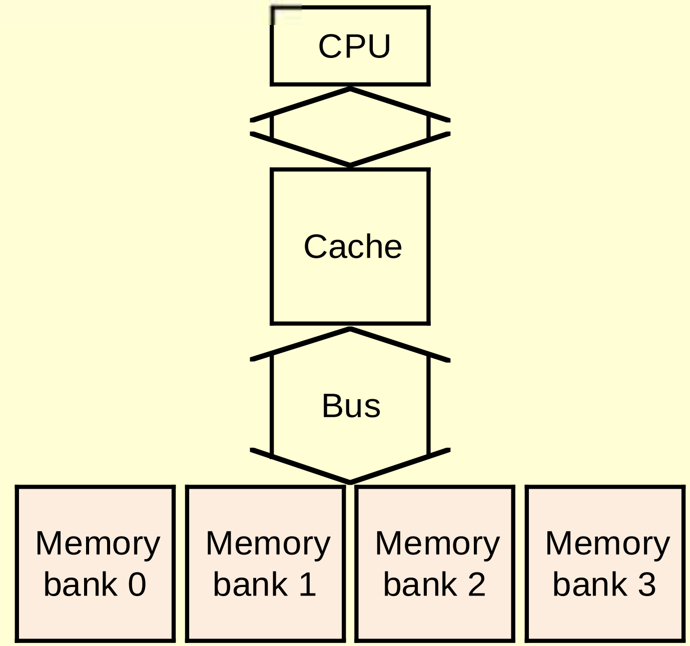

c. Interleaved memory organization

- *Simple*: CPU, Cache, Bus, Memory same width (32 bits)
- *Wide*: CPU/Mux 1 word; Mux/Cache, Bus, Memory N words
- *Interleaved*: CPU, Cache, Bus 1 word: Memory N Modules (4 Modules); example is *word interleaved*

Memory organization would have significant effect on bandwidth

Memory Interleaving

- Access Pattern without Interleaving:

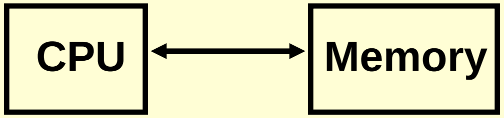

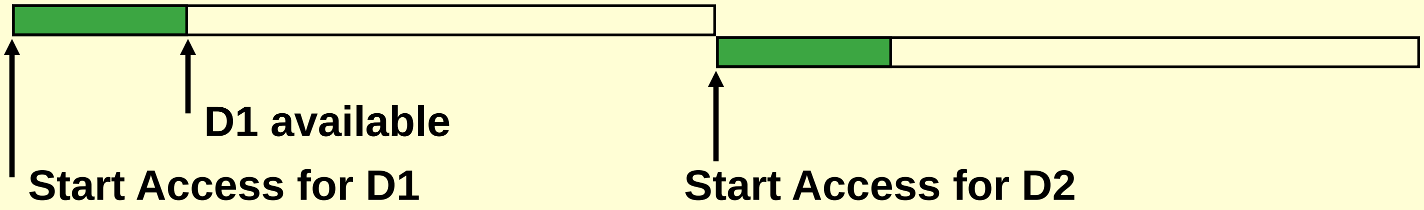

- Access Pattern with 4-way Interleaving:

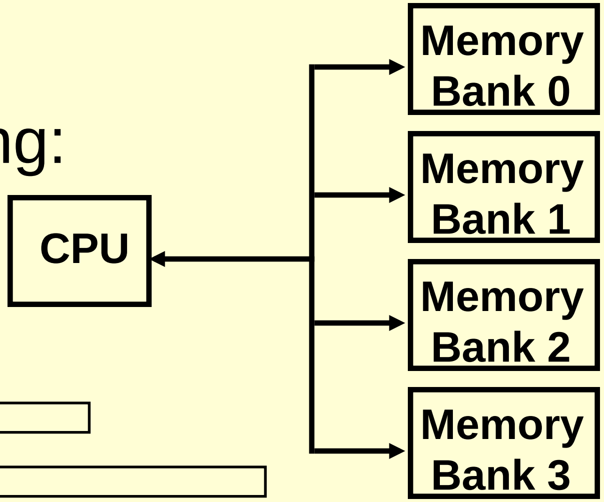

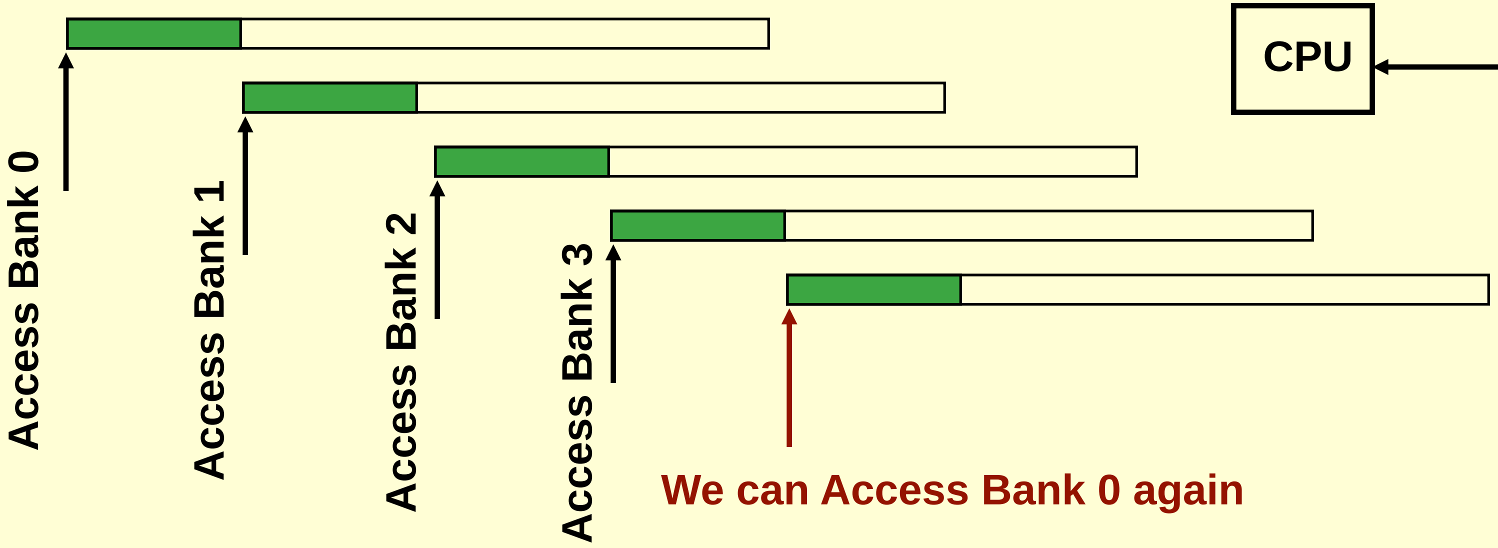

Input/Output

- I/O Interface
  - Device drivers
  - Device controller
  - Service queues
  - Interrupt handling
- Design Issues
  - Performance
  - Expandability
  - Standardization
  - Resilience to failure
- Impact on Tasks
  - Blocking conditions
  - Priority inversion
  - Access ordering

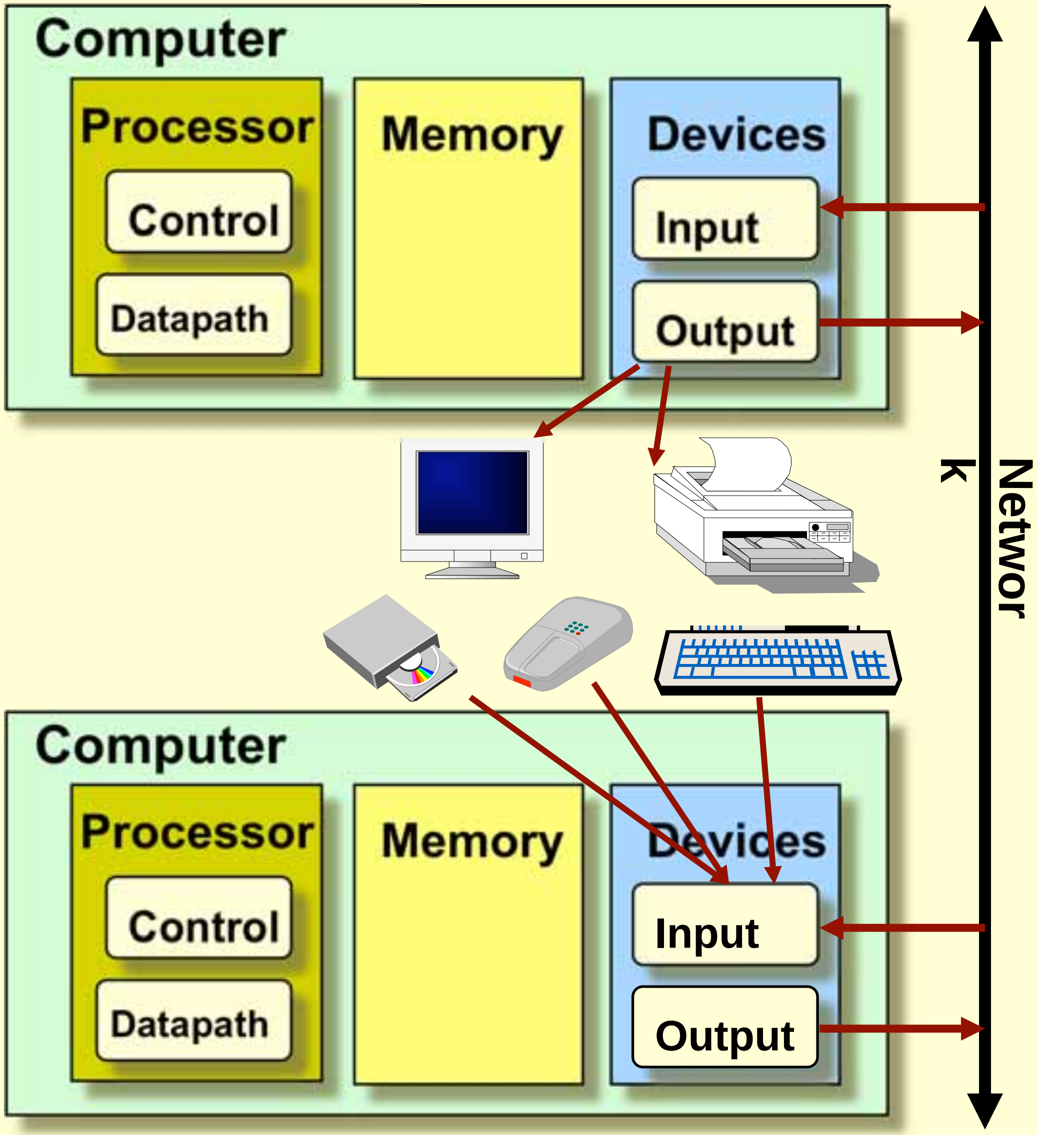

Impact of I/O on System Performance

Suppose we have a benchmark that executes in 100 seconds of elapsed time, where 90 seconds is CPU time and the rest is I/O time. If the CPU time improves by 50% per year for the next five years but I/O time does not improve, how much faster will our program run at the end of the five years?

Answer: Elapsed Time = CPU time + I/O time

After n years	CPU time	I/O time	Elapsed time	% I/O time
0	90 Seconds	10 Seconds	100 Seconds	10%
1	$\frac{90}{1.5} = 60$ Seconds	10 Seconds	70 Seconds	14%
2	$\frac{60}{1.5} = 40$ Seconds	10 Seconds	50 Seconds	20%
3	$\frac{40}{1.5} = 27$ Seconds	10 Seconds	37 Seconds	27%
4	$\frac{27}{1.5} = 18$ Seconds	10 Seconds	28 Seconds	36%
5	$\frac{18}{1.5} = 12$ Seconds	10 Seconds	22 Seconds	45%

Over five years:

CPU improvement = 90/12 = 7. **BUT** System improvement = 100/22 = 4.5

Typical I/O System

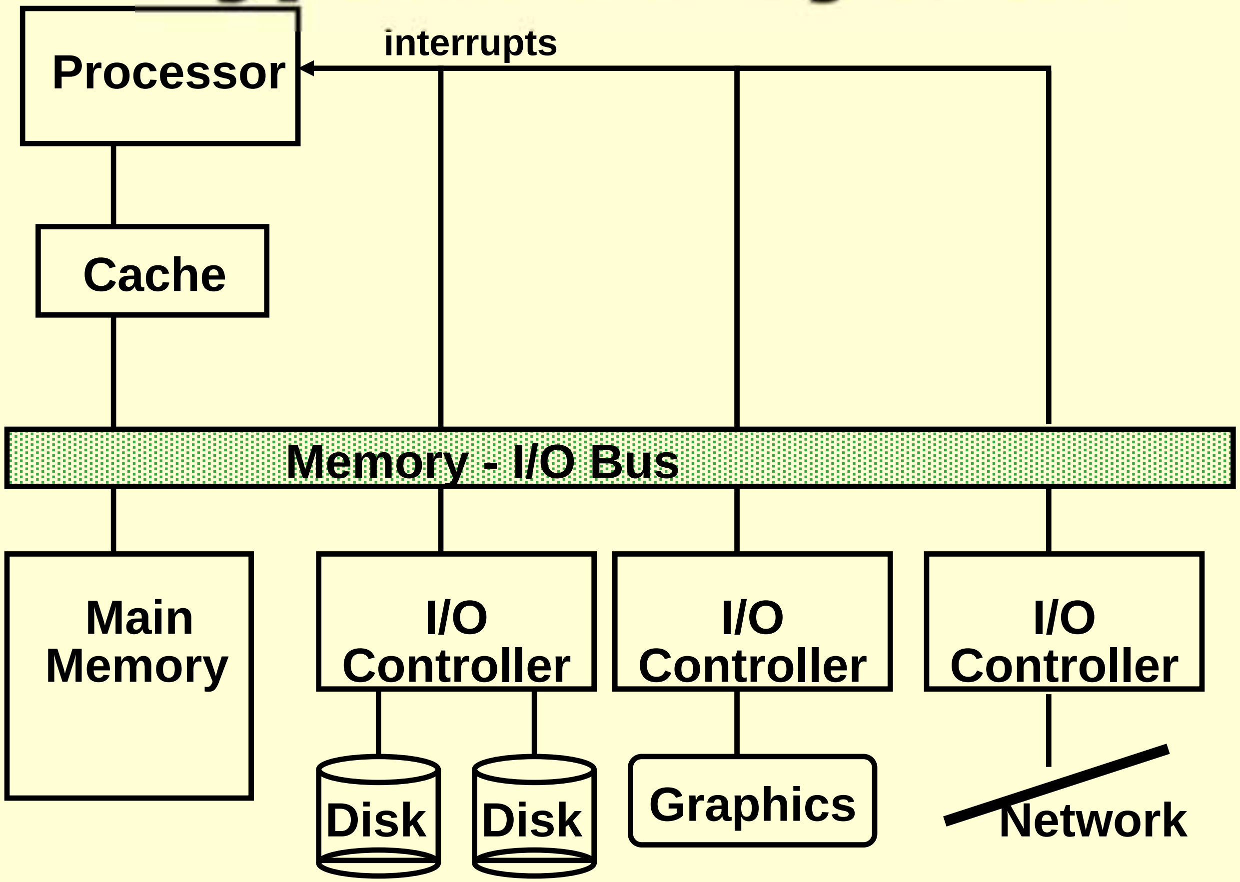

- The connection between the I/O devices, processor, and memory are usually called (local or internal) bus
- Communication among the devices and the processor use both protocols on the bus and interrupts

I/O Device Examples

Device	***Behavior***	***Partner***	***Data Rate (KB/sec)***
Keyboard	Input	Human	0.01
Mouse	Input	Human	0.02
Line Printer	Output	Human	1.00
Floppy disk	Storage	Machine	50.00
Laser Printer	Output	Human	100.00
Optical Disk	Storage	Machine	500.00
Magnetic Disk	Storage	Machine	5,000.00
Network-LAN	Input or Output	Machine	20 – 1,000.00
Graphics Display	Output	Human	30,000.00

Disk History

Data density in Mbit/square inch

Capacity of Unit Shown in Megabytes

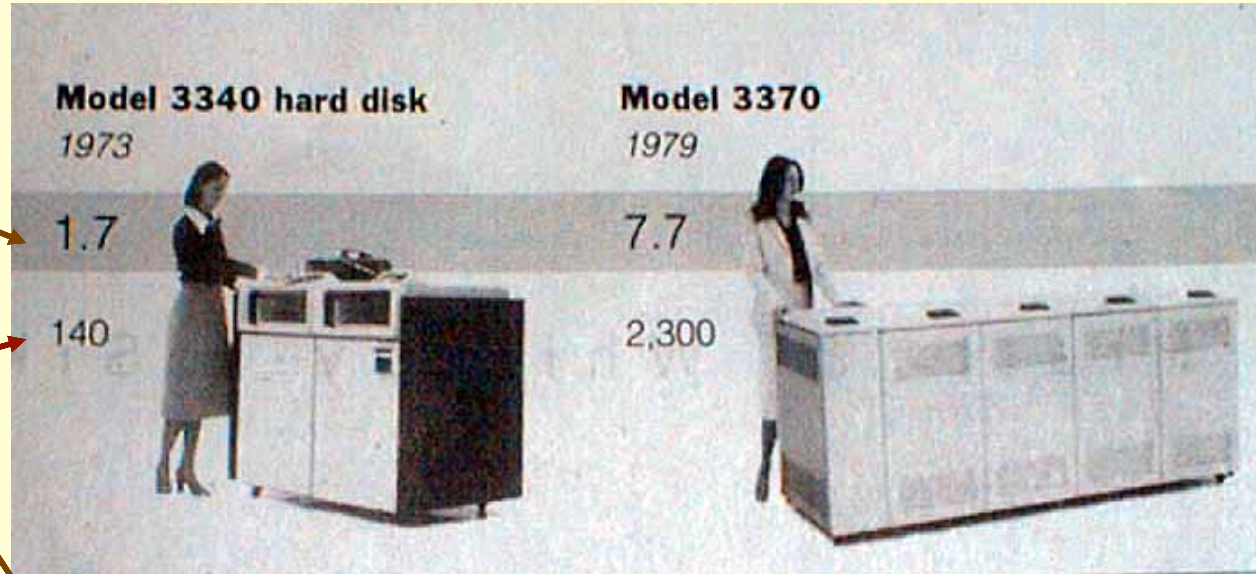

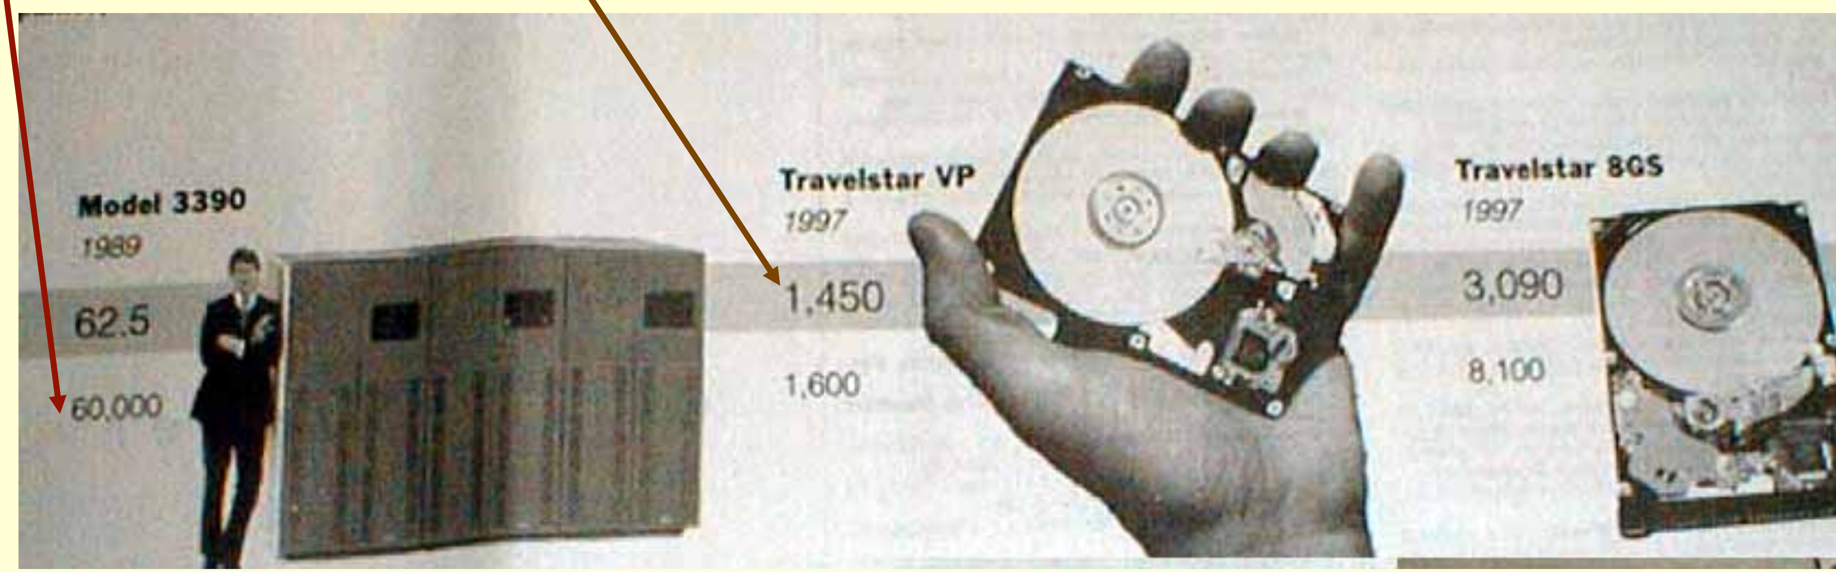

source: New York Times, 2/23/98, page C3

Organization of a Hard Magnetic Disk

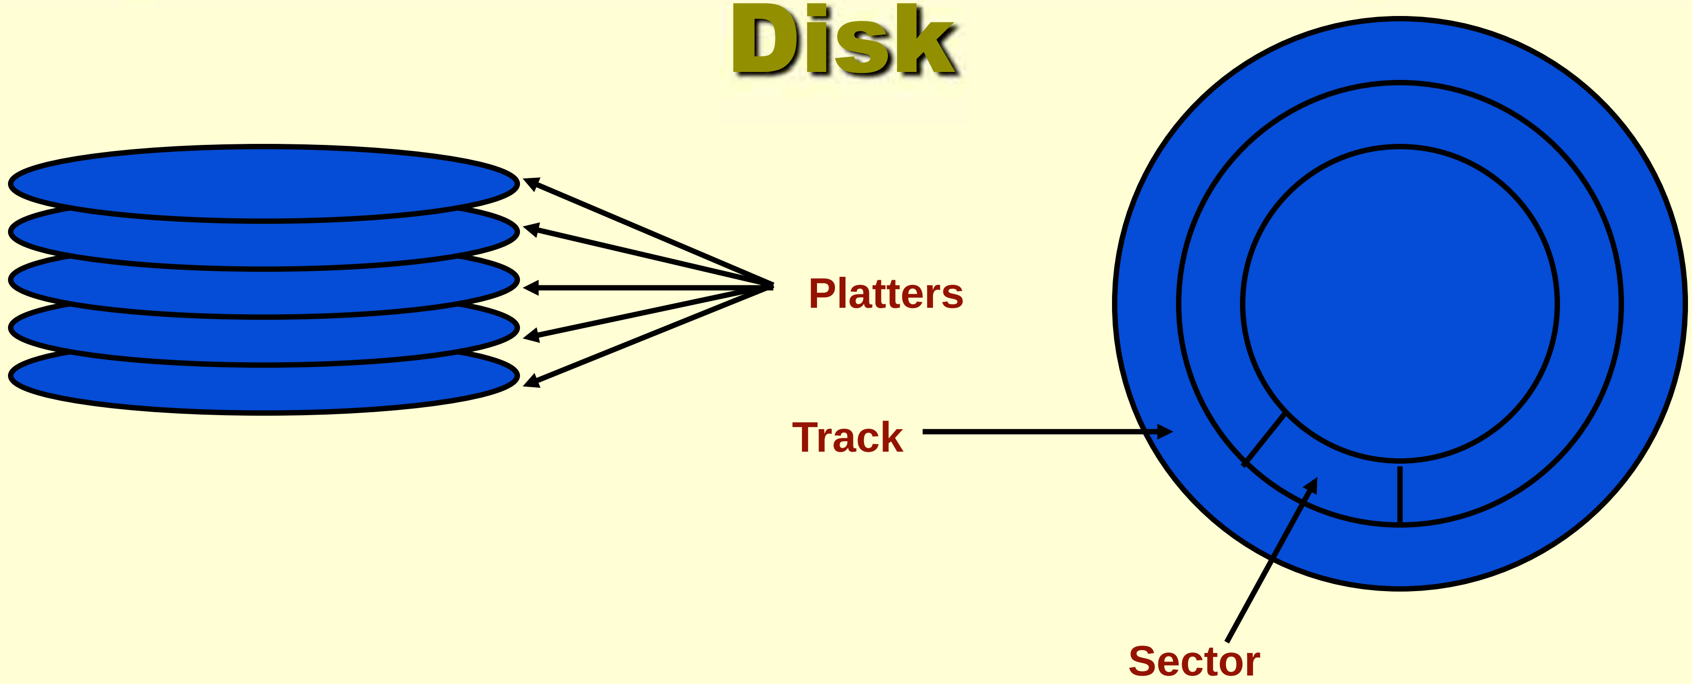

- Typical numbers (depending on the disk size):
  - 500 to 2,000 tracks per surface
  - 32 to 128 sectors per track
    - A sector is the smallest unit that can be read or written to
- Traditionally all tracks have the same number of sectors:
  - Constant bit density: record more sectors on the outer tracks
  - Recently relaxed: constant bit size, speed varies with track location

Magnetic Disk Operation

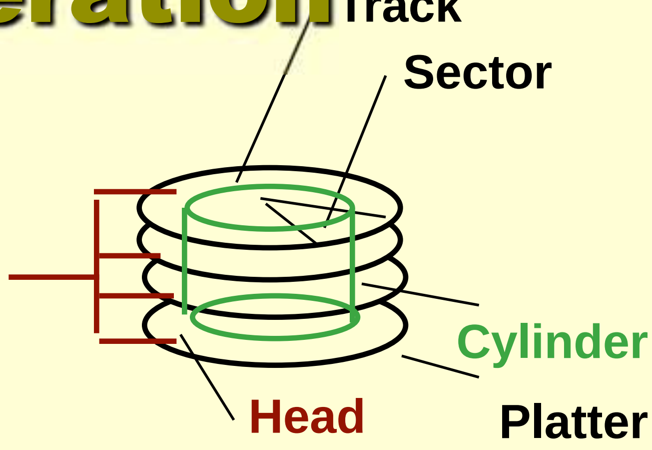

- Cylinder: all the tracks under the head at a given point on all surface
- Read/write is a three-stage process:
  - Seek time
    - position the arm over proper track
  - Rotational latency
    - wait for the sector to rotate under the read/write head
  - Transfer time
    - transfer a block of bits (sector) under the read-write head
- Average seek time
  - (∑ time for all possible seeks) / (# seeks)
  - Typically in the range of 8 ms to 12 ms
  - Due to locality of disk reference, actual average seek time may only be 25% to 33% of the advertised number

Magnetic Disk Characteristic

- Rotational Latency:
  - Most disks rotate at 3,600 to 7,200 RPM
  - Approximately 16 ms to 8 ms per revolution, respectively
  - An average latency to the desired information is halfway around the disk:
    - 8 ms at 3600 RPM, 4 ms at 7200 RPM
- Transfer Time is a function of :
  - Transfer size (usually a sector): 1 KB / sector
  - Rotation speed: 3600 RPM to 7200 RPM
  - Recording density: bits per inch on a track
  - Diameter: typical diameter ranges from 2.5 to 5.25”
  - Typical values: 2 to 12 MB per second

Example

Calculate the access time for a disk with 512 byte/sector and 12 ms advertised seek time. The disk rotates at 5400 RPM and transfers data at a rate of 4MB/sec. The controller overhead is 1 ms. Assume that the queue is idle (so no service time)

Answer:

Disk Access Time = Seek time + Rotational Latency + Transfer time + Controller Time + Queuing Delay

= 12 ms + 0.5 / 5400 RPM + 0.5 KB / 4 MB/s + 1 ms + 0

= 12 ms + 0.5 / 90 RPS + 0.125 / 1024 s + 1 ms + 0

= 12 ms + 5.5 ms + 0.1 ms + 1 ms + 0 ms

= 18.6 ms

If real seeks are 1/3 the advertised seeks, disk access time would be 10.6 ms, with rotation delay contributing 50% of the access time!